

George Karabatsos\* and Stephen G. Walker

# A Bayesian Nonparametric Hypothesis Testing Approach for Regression Discontinuity Designs<sup>1</sup>

February 7, 2014

**Abstract:** The regression discontinuity (RD) design is a popular approach to causal inference in non-randomized studies. This is because it can be used to identify and estimate causal effects under mild conditions. Specifically, for each subject, the RD design assigns a treatment or non-treatment, depending on whether or not an observed value of an assignment variable exceeds a fixed and known cutoff value.

In this paper, we propose a Bayesian nonparametric regression modeling approach to RD designs, which exploits a local randomization feature. In this approach, the assignment variable is treated as a covariate, and a scalar-valued confounding variable is treated as a dependent variable (which may be a multivariate confounder score). Then, over the model's posterior distribution of locally-randomized subjects that cluster around the cutoff of the assignment variable, inference for causal effects are made within this random cluster, via two-group statistical comparisons of treatment outcomes and non-treatment outcomes.

We illustrate the Bayesian nonparametric approach through the analysis of a real educational data set, to investigate the causal link between basic skills and teaching ability.

**Keywords:** Bayesian Nonparametric Regression, Causal Inference, Sharp Regression Discontinuity, Fuzzy Regression Discontinuity.

**Short title:** Bayesian Nonparametric Hypothesis Testing for RD.

---

<sup>1</sup>This research is supported in part by NSF grant SES-1156372.

# 1 Introduction

A basic objective in scientific research is to infer causal effects of a treatment versus non-treatment from empirical data. Randomized studies provide a gold standard of causal inference [19], because such a study ensures, through the random assignments of treatments to subjects, that treatment assignments are independent of all observed and unobserved confounding variables which characterize experimental subjects. As a consequence, the causal effect can be estimated by a direct comparison of treatment outcomes and non-treatment outcomes. However, in many research settings, it is not feasible to randomly assign treatments to subjects; due to financial, ethical, or time constraints [19]. Therefore, in such settings, it becomes necessary to estimate causal effects of treatments from non-randomized, observational designs.

For observational studies, a popular approach to causal inference is based on the regression discontinuity (RD) design ([21], [7]). In the design, each subject is assigned to a treatment or non-treatment whenever an observed value of a continuous-valued assignment variable ( $R$ ) exceeds a cutoff value. The RD design can provide a "locally-randomized experiment", in the sense that the causal effect of treatment outcomes versus non-treatment outcomes can be identified and estimated from subjects having values of the assignment variable that are located in a small neighborhood around the cutoff. As shown in [11], the RD design can empirically produce causal effect estimates that are similar to those from a standard randomized study ([1], [5], [3], [2], [20]). Arguably, the fact that RD designs can provide such locally-randomized experiments is a main reason why, between 1997 through 2013, that at least 74 RD-based empirical studies have emerged in the fields of education, psychology, statistics, economics, political science, criminology, and health sciences ([7], [16], [4], [24], [17]).

Given the locally-randomization feature of RD designs, a conceptually-attractive approach to causal inference is to identify the cluster

$$\mathcal{C}_\varepsilon(r_0) = \{i : |r_i - r_0| < \varepsilon\}$$

of locally-randomized subjects, who have assignment variable observations  $r_i$  that are located in a neighborhood of size  $\varepsilon > 0$  around the cutoff  $r_0$ . Then, for subjects within the cluster  $\mathcal{C}_\varepsilon(r_0)$ , to perform statistical tests by comparing treatment outcomes between subjects with  $r_i \geq r_0$  against non-treatment subjects; i.e. among subjects with  $r_i < r_0$ ) [6].

One approach to identifying the locally-randomized cluster of subjects is to iteratively search for the largest value of  $\varepsilon > 0$  that leads to a non-rejection of the

null hypothesis tests of zero effect of the treatment variable  $T = \mathbf{1}(R \geq r_0)$ . A theoretical motivation for this approach is that, as a consequence of local randomization, the distribution of the confounding variables is the same for non-treatment subjects located just to the left of the cutoff ( $r_0$ ), as for subjects located just to the right of the cutoff [15].

However, this approach is not fully satisfactory. This is because it bases estimates and tests of causal effects (comparisons of treatment outcomes versus non-treatment outcomes), conditionally on an optimal value of the neighborhood size parameter  $\varepsilon$ , i.e., the clustering configuration  $\mathcal{C}_\varepsilon(r_0)$  found via the null hypothesis tests mentioned above. Therefore, the approach does not fully account for the uncertainty that is inherent in the neighborhood size parameter, i.e., the configuration of the cluster of locally-randomized subjects. As a result, the approach may lead to null-hypothesis tests of causal effects that have exaggerated significance levels.

In this manuscript, we propose a Bayesian nonparametric regression approach to estimating and testing for causal effects, for locally-randomized subjects in a RD design, while accounting for the uncertainty in the cluster configuration  $\mathcal{C}_\varepsilon(r_0)$  of locally-randomized subjects. In the approach, we consider a scalar variable  $X$  that we assume sufficiently describes all observed and unobserved pre-treatment "confounding" variables. Again, as a consequence of local-randomization,  $X$  has the same distribution, for subjects located just to the left and for subjects located just to the right of the cutoff, respectively [15].

Then, we fit a Bayesian nonparametric regression model, based on the restricted Dirichlet process (rDP) [23], in order to provide a flexible regression of  $X$  on the assignment variable  $R$ . Importantly, the model provides a random clustering of subjects, with respect to common values of  $X$ , in a way that is sensitive to the ordering of the values of  $R$ . Given a posterior random draw of clusters of subjects under the rDP model, we identify the single cluster

$$\mathcal{C}(r_0) = \{i : \text{individual } i \text{ has an } r_i \text{ in the cluster containing } r_0\}$$

of locally-randomized subjects. Within this cluster we then compare the treatment outcomes (of subjects with  $r_i \geq r_0$ ) against non-treatment outcomes (of subjects with  $r_i < r_0$ ), via statistical summaries and two-sample statistical tests. Specifically, we may compare outcomes in terms of the mean, variance, quantiles, and the interquartile range, and use various two-sample tests, including the  $t$ -test, Wilcoxon-Mann-Whitney test of equality of medians, chi-square test of equality of variances, and the Kolmogorov-Smirnov test. We can even estimate the probability  $\Pr[Y_1 \geq Y_0 | \mathcal{C}(r_0)]$  that the treatment outcome ( $Y_1$ ) exceeds the non-treatment

outcome ( $Y_0$ ) (see [14]). We then average the results of such statistical comparisons over many posterior samples of the clusters  $\mathcal{C}(r_0)$ , under the restricted DP model. Hence, when making such statistical comparisons, we fully account for the uncertainty that is inherent in the clustering configuration  $\mathcal{C}(r_0)$  of locally-randomized subjects. Moreover, this statistical procedure represents another application of inference of posterior functionals in DP-based models [8].

We can extend the Bayesian nonparametric regression procedure to RD design setting that involve multiple pre-treatment confounding variables  $(X_1, \dots, X_p)$ . In this case, we can construct a scalar confounding covariate  $X$  by the multivariate confounder score [18], which enables the comparison of treatment outcomes and non-treatment outcomes, conditionally on subclassified values of the score. Multivariate confounder scores are constructed by a regression of the outcomes  $Y$  on  $(T, X_1, \dots, X_p)$ , and then setting each score to be equal to the part of the predictor that is free of the linear effect of  $T$  on  $Y$ . Moreover, while in this paper we base our causal inference procedure on the restricted DP, in fact, it is possible to base the causal inference procedure on any Bayesian nonparametric regression model that can cluster subjects as a function of the assignment variable  $R$ . See [13], for a recent example.

In the next section, we review the Bayesian nonparametric regression model, based on the rDP. Then we further describe how estimates and tests of the causal effects is undertaken based on the clustering method. We also provide some more details on the multivariate confounder scoring method. Moreover, we show how this Bayesian nonparametric method can be extended to handle causal inferences in a context of a fuzzy RD design [22], which involves imperfect treatment compliance among the subjects. In Section 3, we illustrate our model through the analysis of a data set, to provide causal inferences in an educational research setting. Section 4 ends with conclusions.

## 2 The Bayesian Nonparametric Model

Let  $\{(x_i, r_i, y_i)\}_{i=1}^n$  denote a sample set of data obtained under a RD design. For each subject  $i \in \{1, \dots, n\}$ , the triple  $(x_i, r_i, y_i)$  denotes an observed value of the confounding variable, the assignment variable, and outcome variable, respectively. For such data, we consider a Bayesian nonparametric regression of the pre-treatment confounding variable  $X$  on the assignment variable  $R$ , based on a restricted DP (rDP) mixture of normal linear regressions [23]. Without loss of

generality, we assume

$$r_1 \leq r_2 \leq \dots \leq r_n.$$

The key idea is that clusterings will be based on this order; so the first cluster would be based on the smallest values of  $r$ , and so on. Since here we are interested in clustering subjects on common values of  $X$ , based on values of  $R$ , we represent this model as a random partition model, as follows:

$$\left[ (X_1, \dots, X_n) | r, \rho_n, \{(\beta_j^*, \sigma_j^{2*})\}_{j=1}^{k_n} \right] \sim \prod_{j=1}^{k_n} \prod_{\{i:s_i=j\}} \text{Normal}\left(x_i | \underline{\mathbf{r}}_i^\top \beta_j^*, \sigma_j^{2*}\right) \quad (1a)$$

$$[\rho_n] \sim \pi(\rho_n) = \frac{\alpha^{k_n}}{\alpha^{[n]}} \frac{n}{k_n!} \prod_{j=1}^{k_n} \frac{1}{n_j} \mathbf{1}\{s_1 \leq \dots \leq s_n\} \quad (1b)$$

$$[\beta_j | \sigma_j^2] \sim \text{Normal}(\beta_j | \beta_0, \sigma_j^2 \mathbf{C}^{-1}) \quad (1c)$$

$$[\sigma_j^2] \sim \text{InverseGamma}(\sigma_j^2 | a, b), \quad (1d)$$

where  $\underline{\mathbf{r}}_i = (1, r_i)^\top$ ; the collection  $\{(\beta_j^*, \sigma_j^{2*})\}_{j=1}^{k_n}$  form the  $k_n \leq n$  distinct values in the sample of parameters  $((\beta_1, \sigma_1^2), \dots, (\beta_n, \sigma_n^2))$  that are assigned to each of the  $n$  subjects. And  $\rho_n = (s_1, \dots, s_n)$  is a random partition of the  $n$  observations, where  $s_i = j$  if  $(\beta_i, \sigma_i^2) = (\beta_j^*, \sigma_j^{2*})$ , and with  $n_j = \sum_{i=1}^n \mathbf{1}\{(\beta_i, \sigma_i^2) = (\beta_j^*, \sigma_j^{2*})\}$ .

Also, the rDP is parameterized by a precision parameter  $\alpha$ , and by a normal-inverse-gamma baseline distribution for  $(\beta, \sigma^2)$  that defines the mean of the process. See [23] for more details.

The posterior distribution of the clustering partition  $\rho_n = (s_1, \dots, s_n)$ , with respect to the  $k_n$  distinct values  $\{(\beta_j^*, \sigma_j^{2*})\}_{j=1}^{k_n}$ , is given by:

$$\pi(\rho_n | \mathbf{y}, \mathbf{x}) \propto \frac{\alpha^{k_n}}{k_n!} \prod_{j=1}^{k_n} \frac{1}{n_j} \sqrt{\frac{|\mathbf{C}|}{|\mathbf{C} + \underline{\mathbf{R}}_j^\top \underline{\mathbf{R}}_j|}} \frac{b^a \Gamma(a + n_j/2)}{\Gamma(a)(b + V_j^2/2)^{a+n_j/2}} \mathbf{1}_{s_1 \leq \dots \leq s_n}.$$

In the above,

$$V_j^2 = (\mathbf{y}_j - \hat{\mathbf{y}}_j)^\top \widehat{\mathbf{W}}_j (\mathbf{y}_j - \hat{\mathbf{y}}_j),$$

$$\widehat{\mathbf{W}}_j = \mathbf{I}_j - \mathbf{R}_j^\top (\mathbf{C} + \underline{\mathbf{R}}_j^\top \underline{\mathbf{R}}_j)^{-1} \underline{\mathbf{R}}_j^\top$$

and

$$\hat{\mathbf{y}}_j = \underline{\mathbf{R}}_j \beta_0,$$

with  $\mathbf{R}_j$  the matrix of row vectors  $\mathbf{r}_i = (1, r_i)^\top$  for subjects belonging in cluster group  $j$  [23]. Posterior samples from  $\pi(\rho_n | \mathbf{y}, \mathbf{x})$  can be generated though the use of a reversible-jump Markov Chain Monte Carlo (RJMCMC) sampling algorithm that is described in Section 4 of [23]. At each sampling stage of the algorithm, with equal probability, either two randomly-selected clusters of subjects, that are adjacent with respect to the ordering of  $R$ , are merged into a single cluster; or a randomly selected cluster of subjects is split into two clusters.

Given a random sample of a partition,  $\rho_n \sim \pi(\rho_n | \mathbf{y}, \mathbf{x})$  from the posterior, and given the fixed index  $i \equiv i_0$  of a subject whose assignment variable  $r_i$  is nearest to the cutoff  $r_0$ , we then identify the single (posterior random) cluster of locally-randomized subjects, with this cluster of subjects identified by the subset of indices

$$\mathcal{C}(r_0) = \{i : s_{i_0} = s_i\}.$$

This is a posterior random cluster of subjects with values of the assignment variables  $r_i$  located in a neighborhood around the cutoff  $r_0$ .

For the subjects in this random cluster, we compare treatment outcomes  $y_i$  for subjects where  $r_i \geq r_0$ , versus non-treatment outcomes  $y_i$  for subjects where  $r_i < r_0$ , based on two-sample statistical comparisons of various statistical quantities (e.g., means), as mentioned in Section 1. We then repeat this process over a large number of MCMC samples from the posterior distribution of the partitions,  $\pi(\rho_n | \mathbf{y}, \mathbf{x})$ . We then summarize the posterior distribution of such statistical comparisons over these samples, in order to provide estimates and tests of causal effect of the treatment versus non-treatment, on the outcome  $Y$ .

This procedure can be extended to a fuzzy RD design, where the assignment variable  $R$  represents eligibility to receive a treatment, and some subjects who are assigned treatment  $T_i = \mathbf{1}(r_i \geq r_0)$  opt to receive the other treatment  $(1 - T_i)$ . Then, for the subset of subjects in a given random cluster  $\mathcal{C}_\varepsilon(r_0)$ , we can divide the difference in statistical quantities (e.g., treatment mean minus non-treatment mean) by the difference  $\bar{T}_{\mathcal{C}}^{(r_i \geq r_0)} - \bar{T}_{\mathcal{C}}^{(r_i < r_0)}$ , where  $\bar{T}_{\mathcal{C}}^{(r_i \geq r_0)}$  ( $\bar{T}_{\mathcal{C}}^{(r_i < r_0)}$ , respectively) is the average  $T_i$  for subjects in the locally-randomized cluster  $\mathcal{C}(r_0)$  that have assignment variables  $r_i \geq r_0$  (with assignment variables  $r_i < r_0$ , respectively). Such a divided difference provides an instrumental-variables estimate of a causal effect of treatment versus non-treatment. This is true, provided that the local exclusion restriction holds in the sense that for a given cluster of subjects  $\mathcal{C}_\varepsilon(r_0)$ , any effect of the assignment variable  $\mathbf{1}(R \geq r_0)$  on  $Y$  must be only via  $T$  [16].

For an RD setting that involves multiple pre-treatment confounding variables  $\mathbf{x}_i = (x_{1i}, \dots, x_{pi})^\top$  (for  $i = 1, \dots, n$ ), we construct scalar-valued confounding vari-

able  $x_i$  ( $i = 1, \dots, n$ ) by taking Miettinen's multivariate confounder score [18], with  $x_i = \hat{\beta}_0 + \hat{\beta}_{\mathbf{x}}g(\mathbf{x}_i)$ , based on

$$\hat{E}[Y_i|\mathbf{x}_i, r_i] = \hat{\beta}_0 + \hat{\beta}_{\mathbf{x}}B(\mathbf{x}_i)^\top + \hat{\beta}_R\mathbf{1}(r_i \geq r_0),$$

with  $B(\cdot)$  a chosen (e.g., polynomial) basis transformation of  $\mathbf{x}$ , and with coefficient estimates

$$\hat{\beta} = (\hat{\beta}_0, \hat{\beta}_{\mathbf{x}}, \hat{\beta}_R)$$

obtained by a linear model fit. We consider the Bayesian estimator

$$\hat{\beta} = (v^{-1}\mathbf{I}_q + \mathbf{B}^\top\mathbf{B})^{-1}\mathbf{B}^\top\mathbf{y},$$

with  $\mathbf{y} = (y_1, \dots, y_n)^\top$ , and with  $\mathbf{B}$  the  $(n \times q)$ -dimensional basis matrix with row vectors  $(1, B(\mathbf{x}_i)^\top, \mathbf{1}(r_i \geq r_0))$ ,  $i = 1, \dots, n$ .

### 3 Illustration

A data set was obtained under a partnership between four Chicago University schools of education, which implemented a new curriculum that aims to train and graduate teachers to improve Chicago public school education. This data set involves  $n = 204$  undergraduate teacher education students, each of whom enrolled into one of the four Chicago schools of education during either the year of 2010, 2011, or 2012 (90% female; mean age = 22.5, s.d. = 5.3,  $n = 203$ ); 47%, 21%, 10%, and 22% attended the four universities; 49%, 41% and 10% enrolled in 2010, 2011, and 2012). We investigate the causal effect of basic skills on teacher performance (e.g., [10]), because most U.S. schools of education based their undergraduate admissions decisions on the ability of individual applicants to pass basic skills tests. Here, the assignment variable is a 4-variate random variable, defined by subtest scores on an Illinois test of basic skills, in reading, language, math and writing. Each subtest has a minimum passing score of 240. The dependent variable is the total score on the 50-item Haberman Teacher Pre-screener assessment, and a score in the 40-50 range indicates a very effective teacher. This assessment has a test-retest reliability of .93, and has a 95% accuracy rate in predicting which teachers will stay and succeed in the teaching profession, and is used by many schools to assess applicants of teaching positions [12]. Among all the 205 students of the RD design, the average Haberman Pre-screener score is 29.82 (s.d. = 4.3). The average basic skills score in reading (Read), language

(Lang), math (Math), and writing (Write) was 204.3 (s.d. = 33.4), 204.0 (s.d. = 35.8), 212.6 (s.d. = 42.2), and 238.3 (s.d. = 23.7), respectively.

Using the Bayesian clustering method that was described in the previous section, based on the rDP regression model, we analyzed the data set to estimate the causal effect of passing the reading basic skills exam (treatment), versus not passing (non-treatment), on students' ability to teach in urban schools, for subjects located around the cutoff  $r_0$  of an assignment variable  $R$ . We treated the Haberman z-score as the outcome variable  $Y$ . Also, using the Bayesian rDP regression model, we regressed the scalar-valued confounding variable  $X$  on  $R$ . The confounding variable  $X$  is a multivariate confounder score constructed from 113 pre-treatment variables that describe students' personal background, high school background, and teaching preferences. The multivariate confounder score is based on the the Bayesian linear model fit with prior parameter  $\nu = 1000$ , and with linear basis  $B(\mathbf{x}) = (1, \mathbf{x})^\top$ , as described in the previous section. Also, the assignment variable  $R$  is defined by  $B240d10 = (\min(\text{Read}, \text{Lang}, \text{Math}, \text{Write}) - 240)/10$ . This gives a minimum difference between the four basic subtest scores subtracted by the minimum cutoff score (240), and provides one standard method for handling a multivariate assignment variable [24]. Then, the treatment assignment variable  $T$  is defined by  $\text{BasicPass} = \mathbf{1}(B240d10 \geq 0)$ .

Using code we wrote in the MATLAB (Natick, VA) software language, we analyzed the data set using the Bayesian rDP model. For the model, we chose vague prior specifications  $\beta_0 = \mathbf{0}$ ,  $\mathbf{C} = \text{diag}(10^3, 10)$ ,  $\alpha = a = b = 1$ . All posterior estimates, reported here, are based on 200,000 MCMC samples of clusters  $\mathcal{C}_\varepsilon(r_0)$ , which led to accurate posterior estimates of two-group statistical comparisons, according to standard MCMC convergence assessment criteria [9]. Specifically, univariate trace plots displayed good mixing of model parameters and posterior predictive samples, while all posterior predictive estimates obtained 95% MC confidence intervals with half-width sizes of .00.

Table 1 presents the results of the two-group statistical comparisons, in terms posterior mean and 95% posterior credible interval summaries of the group sample size, mean, variance, interquartile range, skewness, kurtosis, quantiles, the t-statistic, the F-statistic for equality of variances, exceedance probabilities  $\Pr[Y_1 \geq Y_0 | \mathcal{C}_\varepsilon(r_0)]$  and  $\Pr[Y_1 \leq Y_0 | \mathcal{C}_\varepsilon(r_0)]$  of the treatment outcome ( $Y_1$ ) and the non-treatment outcome ( $Y_0$ ), and the Kolmogorov-Smirnov test for the equality of distributions. We find that, in terms of the posterior mean of these statistics, the treatment group in the random clusters  $\mathcal{C}_\varepsilon(r_0)$  (having assignment variable observations  $B240d10 \geq 0$ ) tended to have higher values of the Haberman z-score outcome  $Y$  compared to the non-treatment group in the random clusters  $\mathcal{C}_\varepsilon(r_0)$



(having assignment variable observations  $B240d10 \geq 0$ ), in terms of the mean and quantiles. In contrast, the outcomes of the non-treatment tended to have larger dispersion (variance and interquartile range), and more skewness and kurtosis, compared to the treatment group. The treatment group had a significantly higher 90%ile (.90 quantile) of the z-score outcome  $Y$  compared to the non-treatment group, as the 95% posterior credible intervals of the outcome for the two groups was (1.71, 2.77) and (1.24, 1.71), respectively.

## 4 Discussion

In this paper, we proposed and illustrated a novel, Bayesian nonparametric regression modeling approach to RD designs, which exploits the local randomization feature of RD designs, and which basis causal inferences on comparisons of treatment outcomes and non-treatment outcomes within posterior random clusters of locally-randomized subjects. The approach can be easily extended to fuzzy RD settings, involving treatment non-compliance. We illustrate the Bayesian nonparametric approach through the analysis of a real educational data set, to investigate the causal link between basic skills and teaching ability. Finally, the approach assumes that the RD design provides data on all confounding variables (that are used to construct  $X$ ), but this assumption of no hidden bias is questionable. Therefore, in future research, it would be of interest to extend the procedure, so that it can provide an analysis of the sensitivity of causal inference, with respect to varying degrees of hidden biases, i.e., of effects to hypothetical unobserved confounding variables.

## References

- [1] Aiken, L., S. West, D. Schwalm, J. Carroll, and S. Hsiung (1998): “Comparison of a randomized and two quasi-experimental designs in a single outcome evaluation efficacy of a university-level remedial writing program,” *Evaluation Review*, 22, 207–244.
- [2] Berk, R., G. Barnes, L. Ahlman, and E. Kurtz (2010): “When second best is good enough: A comparison between a true experiment and a regression discontinuity quasi-experiment,” *Journal of Experimental Criminology*, 6, 191–208.

- [3] Black, D., J. Galdo, and J. Smith (2005): “Evaluating the regression discontinuity design using experimental data,” Unpublished manuscript.
- [4] Bloom, H. (2012): “Modern regression discontinuity analysis,” *Journal of Research on Educational Effectiveness*, 5, 43–82.
- [5] Buddelmeyer, H. and E. Skoufias (2004): *An evaluation of the performance of regression discontinuity design on PROGRESA*, World Bank Publications.
- [6] Cattaneo, M., B. Frandsen, and R. Titiunik (2013): “Randomization inference in the regression discontinuity design: An application to the study of party advantages in the U.S. Senate,” Technical report, Department of Statistics, University of Michigan.
- [7] Cook, T. (2008): “Waiting for life to arrive: A history of the regression discontinuity design in psychology, statistics and economics,” *Journal of Econometrics*, 142, 636–654.
- [8] Gelfand, A. and A. Kottas (2002): “A computational approach for full non-parametric Bayesian inference under Dirichlet process mixture models,” *Journal of Computational and Graphical Statistics*, 11, 289–305.
- [9] Geyer, C. (2011): “Introduction to MCMC,” in S. Brooks, A. Gelman, G. Jones, and X. Meng, eds., *Handbook of Markov Chain Monte Carlo*, Boca Raton, FL: CRC, 3–48.
- [10] Gitomer, D., T. Brown, and J. Bonett (2011): “Useful signal or unnecessary obstacle? The role of basic skills tests in teacher preparation,” *Journal of Teacher Education*, 62, 431–445.
- [11] Goldberger, A. (2008/1972): “Selection bias in evaluating treatment effects: Some formal illustrations,” in D. Millimet, J. Smith, and E. Vytlačil, eds., *Modelling and evaluating treatment effects in economics*, Amsterdam: JAI Press, 1–31.
- [12] Haberman, M. (2008): *The Haberman Star Teacher Pre-Screener*, Houston: The Haberman Educational Foundation.
- [13] Karabatsos, G. and S. Walker (2012b): “Adaptive-modal Bayesian nonparametric regression,” *Electronic Journal of Statistics*, 6, 2038–2068.

- [14] Kotz, S., Y. Lumelskii, and M. Pensky (2003): *The Stress-Strength Model and its Generalizations*, New Jersey: World Scientific.
- [15] Lee, D. (2008): “Randomized experiments from non-random selection in U.S. house elections,” *Journal of Econometrics*, 142, 675–697.
- [16] Lee, D. and T. Lemieux (2010): “Regression discontinuity designs in economics,” *The Journal of Economic Literature*, 48, 281–355.
- [17] Li, F., A. Mattei, and F. Mealli (2013): “Bayesian inference for regression discontinuity designs with application to the evaluation of Italian university grants,” Technical report, Department of Statistics, Duke University.
- [18] Miettinen, O. (1976): “Stratification by a multivariate confounder score,” *American Journal of Epidemiology*, 104, 609–620.
- [19] Rubin, D. (2008): “For objective causal inference, design trumps analysis,” *The Annals of Applied Statistics*, 2, 808–840.
- [20] Shadish, W., R. Galindo, V. Wong, P. Steiner, and T. Cook (2011): “A randomized experiment comparing random and cutoff-based assignment,” *Psychological Methods*, 16, 179–191.
- [21] Thistlewaite, D. and D. Campbell (1960): “Regression-discontinuity analysis: An alternative to the ex-post facto experiment,” *Journal of Educational Psychology*, 51, 309–317.
- [22] Trochim, W. (1984): *Research design for program evaluation: The regression discontinuity approach*, Newbury Park, CA: Sage.
- [23] Wade, S., S. Walker, and S. Petrone (2013, to appear): “A predictive study of Dirichlet process mixture models for curve fitting,” *Scandinavian Journal of Statistics*, n/a–n/a.
- [24] Wong, V., P. Steiner, and T. Cook (2013): “Analyzing regression-discontinuity designs with multiple assignment variables: A comparative study of four estimation methods,” *Journal of Educational and Behavioral Statistics*, 38, 107–141.

<b>Statistic</b>	<b>Non-Treatment</b>	<b>Treatment</b>
sample size	103.1 (3, 190)	6.7 (2, 16)
mean	.37 (−.07, 1.55)	1.23 (.97, 1.59)
variance	.76 (.01, 1.04)	.47 (.01, 0.85)
interquartile range	1.17 (.18, 1.71)	.90 (.24, 1.41)
skewness	−.11 (−1.26, .71)	.03 (−.63, 0.82)
kurtosis	2.69 (1.45, 3.41)	2.06 (1.00, 3.06)
1%ile	−1.34 (−2.20, 1.47)	.27 (−0.65, 1.47)
10%ile	−.77 (−1.35, 1.47)	.42 (.01, 1.47)
25%ile	−.20 (−0.65, 1.47)	.74 (.30, 1.47)
50%ile	.34 (−.18, 1.47)	1.22 (1.00, 1.59)
75%ile	.98 (.53, 1.65)	1.65 (1.47, 2.06)
90%ile	1.43 (1.24, 1.71)	2.14 (1.71, 2.77)
99%ile	2.04 (1.71, 2.42)	2.28 (1.71, 2.89)
t-statistic	−2.02 (−4.21, .88)	p-value: .19 (.00, .91)
<i>F</i> test, variance	4.86 (.02, 34.45)	p-value: .65 (.05, .98)
$\Pr[Y_1 \geq Y_0   \mathcal{C}_\varepsilon(r_0)]$	0.70 (.21, .93)	
$\Pr[Y_1 \leq Y_0   \mathcal{C}_\varepsilon(r_0)]$	0.22 (.04, .67)	
KS test	.28 (.05, .98)	

Table 1: Statistical comparisons of treatment outcomes versus non-treatment outcomes.